

# Autonomous Target Tracking of Small Bodies During Flybys \*

Shyam Bhaskaran, Joseph E. Riedel, and Stephen P. Synnott<sup>†</sup>

## Abstract

Spacecraft flybys of small solar system bodies provide important science return in the form of images of the target body taken around closest approach. In order to maximize the number of images taken of the target, an autonomous closed-loop tracking system has been developed to maintain lock on the target during the flyby. The system uses images to estimate the spacecrafts target-relative position and attitude, which is then used to point the camera. The system has been successfully used twice: the Deep Space 1 flyby of comet Borrelly and the Stardust flyby of asteroid Annefrank. This paper describes in detail the tracking algorithms and flight results.

## INTRODUCTION

Within the last decade, a number of deep space missions have flown by small solar system bodies (asteroids and comets). These encounters have either been opportunistic flybys, such as the Galileo spacecrafts encounters with asteroids Gaspra and Ida, or missions specifically targeted for observation of a small body, such as the STARDUST mission's encounter with comet Wild 2. In either case, an important component of the science return from these missions are images taken of the object during the flyby. Because the ephemeris of the small bodies are generally not well determined, a typical imaging sequence during the encounter involves shuttering frames which cover a two to three sigma area of the targets positional uncertainty as projected into the camera field-of-view to guarantee that the target will be in at least one of the frames. Although this process has worked well in the past, it necessarily results in image frames without the target and of no science value.

It is obvious, however, that images taken during the approach to the target provide very good data to improve knowledge of the targets ephemeris beyond what is available from Earth-based observations, but the rapid pace of the flyby coupled with large (in the tens of minutes) round-trip light times preclude processing the observations on the ground to provide this information to the spacecraft. An autonomous onboard system though, can use this data to update its own knowledge of the target ephemeris and execute an imaging sequence taking advantage of this information to dramatically increase the number of image frames which include the target, thereby increasing the science return. Such a system has been developed for use, extensively tested on the ground using simulations, and proven in flight on three occasions – the Deep Space 1 (DS1) flyby of comet Borrelly and the STARDUST flybys of asteroid Annefrank and comet Wild 2. Because the tracking system uses a reduced set of dynamics as compared to a standard orbit determination setup, the system was dubbed RSEN, for Reduced State Encounter Navigation. This paper describes the algorithms used by RSEN, as well as results from the ground testing and flight results.

---

\*Copyright (c) 2004 by the California Institute of Technology. Permission to publish granted to The American Astronautical Society.

<sup>†</sup>Navigation and Flight Mechanics Section, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California, MS 264-380, 4800 Oak Grove Dr., Pasadena, CA 91109, Ph: (818)354-3152, Email: Shyam.Bhaskaran@jpl.nasa.gov

## BACKGROUND

Before describing the details of the tracking algorithm, a brief mention should be made of the context in which has been used. This section will provide information on the missions which have employed RSEN, the navigation leading up to the closed-loop tracking phase, the camera setup, and some design considerations.

### The Missions

The first mission to image a small body was the European Giotto spacecraft which flew by comet Halley in 1986. Although Giotto had an autonomous onboard system to track Halley's nucleus, the camera failed, probably due to impacts with cometary dust, 50 seconds before closest approach. Nevertheless, images on approach were available and revealed the nucleus to resolutions on the order of several 100 m. Subsequently, the Galileo spacecraft's trajectory had opportunistic encounters with the asteroids Gaspra and Ida in 1991 and 1993 respectively. Using a mosaic strategy, whereby the image frames are laid out in a two-dimensional pattern to cover uncertainties in the position of the asteroid relative to the spacecraft, high resolution images were obtained of these asteroids. A similar strategy was also used by the NEAR spacecraft's flyby of asteroid Mathilde prior to its eventual orbit around asteroid Eros.

Deep Space 1 was primarily a technology validation mission. It supported 12 new technologies, including the first ion propulsion engine for deep space, advanced solar arrays, and an autonomous navigation system (autonav). The mission was launched in October 1998 and flew by the asteroid Braille in July 1999. This phase completed its prime requirement of validating the technologies; an extended mission was approved to encounter the short period comet Borrelly in 2001. Autonav was successfully demonstrated for use in interplanetary cruise during the prime mission, but RSEN, which was a subset of autonav, failed to track Braille during the flyby. The primary cause of this failure was the fact that an experimental channel on the camera used to image the asteroid had insufficient sensitivity to detect a signal from the asteroid at the exposure times used. Without a detectable signal from the asteroid, RSEN was unable to lock onto it and no usable science images were obtained. Subsequently, improvements were made to the RSEN algorithm to make it more robust, and the experimental camera was dropped in favor of a standard Charge-Coupled-Device (CCD) for use on the Borrelly flyby. On September 22, 2001, DS1 flew by Borrelly and successfully imaged the nucleus at a resolution of about 45 m.

STARDUST is the first mission by NASA dedicated to a comet. Its prime goal is to collect dust samples from the coma of comet Wild 2 and return them to Earth. During the flyby, a secondary goal is to image the nucleus. The spacecraft was launched on February 1999, flew by asteroid Annefrank on November 2002 and comet Wild 2 on January 2004, and is scheduled to return to Earth on January 2006. The main purpose of the Annefrank encounter was to test the procedures that were to be used on Wild 2, in particular, the autonomous tracking system. The RSEN algorithm onboard STARDUST is very similar to the one on DS1 and successfully tracked both Annefrank and Wild 2. The latter encounter produced images with a resolution of about 14 m, the highest of a comet to date.

### Navigation

The primary mode of navigation for deep space missions is based on radio tracking data obtained through one of the Deep Space Network antenna complexes located in California, Spain, and Australia. Two standard data types are employed: two-way coherent Doppler, which measures the line-of-sight velocity of the spacecraft relative to the station, and two-way coherent range, which measures the line-of-sight distance. The observed data are differenced with predicted measurements based on a mathematical model of the spacecraft's trajectory, taking into account the various gravitational and non-gravitational forces which act on it, and a least-squares fit to the trajectory is obtained. The radio data alone is sufficient for many classes of deep space missions, especially those to the inner planets whose orbits are fairly well known.

For small bodies, however, radio data alone is not adequate because the body's ephemerides are not as accurately known. In order to image an asteroid or comet during a flyby, the spacecraft's target relative position must be known to better than 10 km. Typical ground-based ephemeris knowledge of asteroids ranges from the several tens of km for the larger, brighter ones, to over a 100 km for smaller ones. Comets, due to the influence of non-gravitational forces acting on them from jets and other outgassing sources, have orbit uncertainties on the order of several thousand km. Since the gravitational attraction on the spacecraft due to the body is negligible, the radio data obtained on approach cannot determine where the target is. Thus, optical images taken from the spacecraft are used to improve the ground-based target ephemeris and simultaneously solve for the spacecraft's state relative to the target.

Depending on the characteristics of the camera and the geometry of the approach, a standard imaging sequence starts weeks or months prior to the encounter. The shuttered images are sent to the ground, where they are processed along with the radio data. Since the image data is effectively an angular measurement, the target relative knowledge of the spacecraft state improves as it gets closer to the body. Due to the rapid pace of events immediately surrounding the encounter and the round-trip light times, the ground in the loop image processing is halted 12-24 hours prior to encounter, and the sequences for imaging built using the best available knowledge to that point. This is good enough to permit a mosaic sequence, which covers the 3 sigma uncertainty in the state knowledge, using 10-20 images in the mosaic with the expectation that the target will be in several of the frames. As described earlier, the Galileo and NEAR spacecraft used this technique to obtain close encounter images of their respective targets.

The motivation for developing RSEN was that improvements in the accuracy of the spacecraft's orbit is available in the images taken near approach and could be utilized to maximize the imaging science return. Since it was impractical to process the images on the ground, however, the processing would have to be done onboard. Because DS1's main purpose was to demonstrate new technologies, it was the perfect mission to prove the new method.

## Imaging System

Before describing the tracking algorithm, a brief description should be made of the imaging systems used. Although the descriptions here are specific to the two missions (DS1 and STARDUST) which have employed RSEN, the algorithms themselves are not dependent on a particular set of imaging hardware.

The camera used on DS1 was part of a multi-instrument package called MICAS (Miniature Integrated Camera and Spectrometer) which was one of the experiments flown onboard the spacecraft. The camera had a 600 mm focal length lens with a 1024x1024 pixel array Charge-Coupled-Device (CCD) to record the image. The resultant field of view (FOV) was roughly 13 microradians per pixel, for a total FOV of 1.3 milliradians, or 0.77 deg. The camera was hard-mounted on the spacecraft bus; to point the camera in a particular direction, the entire spacecraft had to be slewed to the proper attitude. The digitization on the images was 12 bit, so brightness value data numbers (DN) ranged between 0 (black) to 4095 (white).

For STARDUST, the camera had a focal length of 201 mm and focused light to a 1024x1024 pixel array CCD. The pixel FOV was 60 microradians, and the total FOV was 61 millirad, or about 3.5 deg. Although the hardware was capable of 12 bit digitization, for increased throughput and storage during encounter, the DN values were square root compressed to 8 bit, resulting in DN values between 0 and 255. As for DS1, the camera itself was fixed to the spacecraft, however, the light path also included a scan mirror which could rotate about a single axis over a range of about 200 deg. Thus, the camera boresight could cover a 3.5 deg by 200 deg range without having to reorient the spacecraft.

The inclusion of the scan mirror was important for the comet flyby as the spacecraft had to keep the main bus behind a dust shield to protect against particle impacts. The dust shield is located on the +X axis of the spacecraft, and so the only attitude changes allowed were rotations about this axis. The scan mirror swept the FOV in the spacecraft X-Z plane, but the *a priori* knowledge of

the spacecraft’s ephemeris relative to the target at the initiation of RSEN was not good enough to know exactly what the orientation of the plane should be. Thus, the encounter sequence included a planned roll about the  $X$  axis six minutes prior to encounter to align the plane using the improved trajectory information from RSEN.

The remainder of this paper will describe the RSEN system used by DS1 and STARDUST, with the focus on the STARDUST specific version. Although the basic algorithm is identical, some differences existed due to differences in spacecraft and mission characteristics; these differences will be noted where appropriate.

## RSEN BASICS

### RSEN Design Considerations

Key considerations in the design of RSEN were that it must be simple, robust, and fast. Computing resources onboard spacecraft generally lag several years behind that available on the ground, and furthermore, many systems onboard that are required for spacecraft health and safety are competing for time on the same processor. This necessitates keeping the algorithm and associated code fairly simple, which also helps lower processing speed. Also, unlike the case of other optical tracking systems such as missile interceptors, the imaging frames used for tracking are taken at relatively infrequent intervals: the maximum rates are on the order of one every 5 to 10 seconds. Thus, a scheme whereby the target is tracked by simply “following the bright spot” will not be as reliable since in the span of the images, spacecraft motions due to attitude deadband excursions will not always be consistent. To maximize the robustness, the best approach is to solve the fundamental problem, that is, improve the knowledge of the spacecraft’s target relative orbit in order to guarantee that the target is tracked. A problem that is generally formulated in terms of attitude control now becomes an exercise in orbit determination, and is solved through a Kalman filtering process to compute an orbit from observations. The steps in this process will now be described.

### Observable Generation

The sole data type used by the filter are images taken of the target body from the spacecraft’s camera. By determining the center of the nucleus in the image, the line-of-sight (LOS) direction of the comet from the spacecraft can be computed. Explicitly though, the LOS to the nucleus is not actually computed from the image; instead, the pixel and line (the  $x$  and  $y$  coordinates in the CCD image) location of the nucleus center is determined from the image and differenced with predicts of the center location to obtain data residuals. In order to compute predicts of the comet location in the camera FOV, the transformation of an inertial vector into camera pixel and line coordinates is needed. This is a three step process; the first step is to rotate an inertial vector into a camera coordinate frame (the  $M$ - $N$ - $L$  frame shown in Figure1), the second is to project these 3-D coordinates into the 2-D camera focal plane, and then finally scale the result into values of pixel and line.

The first step requires the inertial to spacecraft body-fixed rotation matrix,  $\mathbf{T}_{IBF}$ . This is provided by the spacecraft’s Attitude Control System (ACS) using information from the star tracker or gyroscopes. Then, the rotation to the camera  $M$ - $N$ - $L$  coordinate system  $\mathbf{T}_0$  is needed. This rotation can be fairly simple, as for DS1 where the camera was hard-mounted on the spacecraft bus, or somewhat lengthy in the case of STARDUST where the light path included a swiveling mirror. Details of this computation can be found elsewhere<sup>1</sup>; here it suffices to lump it into one rotation matrix,  $\mathbf{T}_0$ , which pre-multiplies  $\mathbf{T}_{IBF}$ . The transformation from inertial to camera frame,  $\mathbf{T}_{IC}$  is then

$$\mathbf{T}_{IC} = \mathbf{T}_0 \mathbf{T}_{IBF}. \quad (1)$$




Figure 1: Camera Focal Plane Coordinates

An inertial LOS vector,  $\mathbf{V}_I$ , can then be rotated into a vector in the camera coordinates,  $\mathbf{V}_C$  by

$$\mathbf{V}_c = \begin{bmatrix} V_{c1} \\ V_{c2} \\ V_{c3} \end{bmatrix} = \mathbf{T}_{IC} \mathbf{V}_I. \quad (2)$$

Once  $\mathbf{V}_c$ , a LOS vector in camera  $M$ - $N$ - $L$  coordinates is obtained, it needs to be transformed into the 2-D camera focal plane. A detailed description of this process can be found elsewhere<sup>2</sup>; a brief synopsis will be given here. First, apply the gnomonic projection,

$$\begin{bmatrix} x \\ y \end{bmatrix} = \frac{f}{V_{c3}} \begin{bmatrix} V_{c1} \\ V_{c2} \end{bmatrix} \quad (3)$$

where

- $f$  = the camera focal length, in mm
- $V_{c1}, V_{c2}, V_{c3}$  = the components of the line-of-sight vector in  $M$ - $N$ - $L$  coordinates
- $x, y$  = the projection of the LOS vector into focal plane coordinates, measured in mm.

Next, find the bias to  $x$  and  $y$ ,  $\Delta x$  and  $\Delta y$ , caused by optical distortions by:

$$\begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} -yr & xr^2 & -yr^3 & xr^4 & xy & x^2 \\ xr & yr^2 & xr^3 & yr^4 & y^2 & xy \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \\ \nu_4 \\ \nu_5 \\ \nu_6 \end{bmatrix} \quad (4)$$

where  $r = x^2 + y^2$ , and the  $\nu$ 's are the optical distortion coefficients. The corrected image locations,  $x'$  and  $y'$ , are then

$$\begin{bmatrix} x' \\ y' \end{bmatrix} = \begin{bmatrix} x + \Delta x \\ y + \Delta y \end{bmatrix}. \quad (5)$$

Finally, the conversion from the rectangular coordinates to pixel and line is:

$$\begin{bmatrix} p \\ l \end{bmatrix} = \begin{bmatrix} K_x & K_{xy} & K_{xxy} \\ K_{yx} & K_y & K_{yyx} \end{bmatrix} \begin{bmatrix} x' \\ y' \\ x'y' \end{bmatrix} + \begin{bmatrix} p_o \\ l_o \end{bmatrix}, \quad (6)$$

where the elements of  $\mathbf{K}$  is a transformation matrix from mm to pixel/line space, and  $p_o$  and  $l_o$  are the center pixel and line of the CCD. Exact values for all the camera parameters are obtained in flight by taking images of dense star fields. Using the known positions of the stars from an accurate star catalog, the observed locations are compared against predicted ones and the parameters are adjusted in a least-squares fit. With the calibrations, the precision of an optical observable will generally be in the range of 0.1 to 0.2 pixels.

## Centerfinding

Centerfinding is the process of obtaining the center of the target object in the image for use by the filter. Historically, this process has been refined through the Voyager encounters with Uranus and Neptune such that accuracies of much less than a pixel are possible.<sup>3</sup> The pictures taken by the spacecraft of planetary satellites were sent to the ground for processing. An analyst would first determine a rough center by eye; computer algorithms which model the ellipsoidal shape of the target with correct lighting due to phase angles and albedo variations could then refine the center location guess to a high accuracy. For autonomous tracking of small bodies however, the centroiding must be done onboard without human intervention on an object with unknown size, shape and albedo properties, and possibly even outgassing in the case of comets. Also, stray light and bright spots due to “hot” pixels or cosmic rays must also be accounted for.

For these reasons, a simple moment algorithm to compute a center-of-brightness was found to lack the robustness required. Instead, a technique termed the “blobber” was developed.<sup>4,5</sup> The blobber searches the image and finds all bright pixels which pass a threshold criteria, both on the low and high end. Contiguous pixels above the threshold are lumped together to form blobs, or regions where the object being searched for could possibly be. The blobs are then sorted in terms of size, and the largest blob which meets a given size criteria is assumed to be the target. In this manner, small spots such as those due to cosmic rays, or large swaths of stray light, can be eliminated as candidate targets. The pixel and line values of the center-of-brightness inside the candidate blob is computed using a simple moment algorithm, and this becomes the observable.

In principle, the desired observable is the center-of-mass of the target since the brightness center can change due to viewing and lighting geometries. Thus, experiments were tried where an empirical shift was applied to the center-of-brightness to approximate the offset from the mass center. Due to lack of knowledge of the true body shape and mass properties, however, the offset was found in practice to offer no real improvements in accuracy and was abandoned.

## Dynamics Model

The time frame during which the onboard navigation processing will be active is about the 30-40 minutes surrounding closest approach. Because the flyby distance is fairly large ( $> 100$  km) and the mass of the target body is fairly small, the trajectory during this time is essentially linear (perturbations caused by impacts with comet dust particles is also negligible). Thus, the trajectory model used by the filter can safely assumed to be a straight line,

$$\mathbf{r}(\mathbf{t}) = \mathbf{r}(\mathbf{t}_0) + \dot{\mathbf{r}}(\mathbf{t} - \mathbf{t}_0) \quad (7)$$

$$\dot{\mathbf{r}}(\mathbf{t}) = \dot{\mathbf{r}}(\mathbf{t}_0). \quad (8)$$

where  $\mathbf{r}$  is the three-dimensional cartesian position and  $\dot{\mathbf{r}}$  is the velocity, expressed in the body-centered, J2000 Earth Mean Equatorial inertial coordinate system. The initial conditions,  $\mathbf{r}(\mathbf{t}_0)$  and  $\dot{\mathbf{r}}(\mathbf{t}_0)$  when onboard navigation is started are provided by the results from ground-based navigation. The uncertainties in these initial conditions is largely in the position, whereas the velocity is well determined from Doppler data to better than 10 cm/s. For this reason, the onboard filter only needs to update the position; the velocity is assumed to be perfect and is not updated. Thus, corrections to the target-centered position,  $\Delta\mathbf{r} = [\Delta\mathbf{x} \ \Delta\mathbf{y} \ \Delta\mathbf{z}]^T$ , form the first three components of the estimate vector.

It was recognized fairly early that the knowledge of the spacecraft attitude was a major error source during encounter. Attitude is generally determined by the use of star trackers, which provide absolute inertial reference without biases or drifts; errors in the estimated attitude are random and would cause no problem for RSEN. For entirely different reasons, however, star trackers were not available for either DS1 or STARDUST during their respective encounters. On DS1, the single star tracker onboard failed shortly after the Braille encounter. Although the remainder of the cruise phase of the mission could use the MICAS CCD as a substitute star tracker, the encounter had to rely on the gyroscopes for attitude knowledge. Without proper calibration from star trackers, the gyros had substantial amounts of drift over short periods of time. If not properly modeled, the gyro drift could easily be mistaken for translational spacecraft motion and result in loss of lock during tracking. The characteristics of the drift were not strictly linear, but over the 30 minute interval RSEN would be operating, a linear approximation was sufficient. Furthermore, because the camera was hard-mounted to the spacecraft bus, the drifts mapped directly into camera pixel and line. Thus, the DS1 flavor of RSEN also estimated, in addition to the three components of position, four additional parameters:  $p_d$  and  $l_d$ , the pixel and line initial gyro bias value, and  $\dot{p}_d$  and  $\dot{l}_d$ , the gyro drift rate.

On STARDUST, there was a concern that the star tracker would lose lock through encounter, either due to coma opacity causing stars to not be visible, or dust particles confusing the star pattern matching. Thus, about five hours prior to encounter, the star tracker was turned off and attitude determined solely by gyros. The drift rates for the STARDUST gyros were small enough to not pose a problem for RSEN, but the initial attitude bias when the attitude estimation was switched to gyros was. In particular, due to the rotating mirror system, the initial offset in attitude causes a signature in the pixel/line residuals which once again, could be mistaken by the filter as translational errors unless modeled. Also due to the rotating mirror system, attitude biases in all three axes had to be modeled since through the course of the flyby, errors in each axis will project into the camera FOV. Since the signature of the error in the FOV is not a simple linear function as it was for DS1, it could not be modeled in pixel/line space; instead, the attitude offset itself is estimated. For the STARDUST flavor of RSEN, the filter estimates, in addition to position, corrections to three components of spacecraft attitude: the right ascension ( $\lambda$ ) and declination ( $\delta$ ) of the body-fixed spacecraft  $X$ -axis, and twist ( $\varphi$ ), the rotation about this axis. Note that these parameters are directly related to the attitude rotation matrix,  $\mathbf{T}_{\mathbf{IBF}}$  from Eq. 1, by:

$$\mathbf{T}_{\mathbf{IBF}} = \mathbf{R}_1(\varphi)\mathbf{R}_2(-\delta)\mathbf{R}_1(\lambda), \quad (9)$$

where  $\mathbf{R}_1$  and  $\mathbf{R}_2$  are Euler angle rotations about the spacecraft  $X$  and  $Y$  axes, respectively.

## Filter Equations

The filter used to provide a state update at the current time,  $t_i$  is a standard extended Kalman filter. Due to the fact that the translational equations of motions are linear, and only corrections to the nominal attitude are needed, numerical integration is not needed which greatly simplifies the filter. The relationship between the observables (the pixel and line coordinates of the center-of-brightness,  $p$  and  $l$ ), and the spacecraft state and attitude over the time span,  $t$ , is:

$$p = p(\mathbf{r}, \lambda, \delta, \varphi, \mathbf{t}) + p_d + \dot{p}_d t, \quad (10)$$

$$l = l(\mathbf{r}, \lambda, \delta, \varphi, \mathbf{t}) + l_d + \dot{l}_d t. \quad (11)$$

The functional dependence of  $p$  and  $l$  on the position and attitude are given by Eqs. 1 through 6 and Eq. 9. The term in brackets only applies to DS1 to estimate the gyro bias and drift. The estimated parameters,  $\mathbf{X}$  for DS1 is

$$\mathbf{X} = [\Delta x \ \Delta y \ \Delta z \ p_d \ \dot{p}_d \ l_d \ \dot{l}_d] \quad (12)$$

and for STARDUST,

$$\mathbf{X} = [\Delta x \ \Delta y \ \Delta z \ \Delta \lambda \ \Delta \delta \ \Delta \varphi]. \quad (13)$$

The partial derivatives of the observables with respect to the state,  $\mathbf{H}$  are:

$$\mathbf{H} = \begin{bmatrix} \partial p / \partial \mathbf{X} \\ \partial l / \partial \mathbf{X} \end{bmatrix} \quad (14)$$

$\mathbf{H}$  is thus a  $2 \times 7$  array for DS1 and a  $2 \times 6$  array for STARDUST. The derivation of the partial derivatives of pixel and line with respect to position and attitude can be found elsewhere<sup>2</sup>; the partials with respect to the pixel and line drifts for DS1 are simply 1 and  $t$ .

Since the only estimated parameters are the position and constant attitudes, the state transition matrix to map estimates from  $t_{i-1}$  to  $t_i$  is the identity matrix. Thus, the a-priori covariance at  $t_i$  is the same as for  $t_{i-1}$ , that is,

$$\bar{\mathbf{P}}_i = \bar{\mathbf{P}}_{i-1}. \quad (15)$$

The standard form of the Kalman gain matrix can then be written as:

$$\mathbf{K}_i = \bar{\mathbf{P}}_i \mathbf{H}_i^T (\mathbf{H}_i \bar{\mathbf{P}}_i \mathbf{H}_i^T + \mathbf{R}_i)^{-1}. \quad (16)$$

Here,  $\mathbf{R}_i$  is a diagonal weighting matrix,

$$\mathbf{R}_i = \begin{bmatrix} \sigma_p^2 & 0 \\ 0 & \sigma_l^2 \end{bmatrix}, \quad (17)$$

where  $\sigma_p^2$  and  $\sigma_l^2$  are the weights on pixel and line, respectively. The data weights were computed according to the following formula:

$$\sigma(p, l)^2 = A_n^2 + (k S_c)^2. \quad (18)$$

Here,  $A_n$  is a component to account for the jitter in the camera due to random attitude noise, while  $S_c$  is the size of the object as projected into the camera focal plane, measured in pixels, with a scale factor  $k$ . The values of  $A_n$  and  $k$  were parameters which could be tuned experimentally using Monte Carlo simulations to see which combinations gave the best performance. The final values chosen were 50 and 0.05 for  $A_n$  and  $k$  respectively. The updated estimate of the state,  $\hat{x}$ , is then given by:

$$\hat{x} = \Delta \mathbf{X} = \mathbf{K}_i \begin{bmatrix} p_o - p_c \\ l_o - l_c \end{bmatrix}, \quad (19)$$

where the observed centers,  $p_o$  and  $l_o$  are obtained from the centroiding process, and  $p_c$  and  $l_c$  are computed using Eqs. 1-6 and the nominal value for  $\mathbf{X}$ . Finally, the updated covariance at  $t_i$  is calculated as:

$$\mathbf{P}_i = (\mathbf{I} - \mathbf{K}_i \mathbf{H}_i) \bar{\mathbf{P}}_i. \quad (20)$$

Given the equations for the Kalman filter, the centroiding process, and the inertial to camera transformations, the algorithm to do the updates can be described. Prior to starting the autonomous tracking, the software is initialized with the current camera model, predicted nucleus size, and the spacecraft state (position and velocity) as determined from ground-based navigation at the start time. The covariance on the initial position is also provided. The starting time is chosen to be early enough that the target is guaranteed to be in the camera FOV when the first image is shuttered, but late enough such that the body is of sufficient size to distinguish it from random spikes. With these constraints, 20-30 minutes prior to closest approach is a good start time. As each image is taken by the camera (at a nominal frequency of one every 10-30 seconds), the tracking software is run to determine object centroids, which are then stored. Along with each image, information about the shutter time and spacecraft attitude as determined from the gyros is passed along. After 15-20 minutes when a sufficient number of images have been processed, the residuals are passed through a data editor to remove outlier points. The data editor compares three contiguous points in a moving window; if any point is not consistent with its neighbors, it is deleted. Once this step is completed, the filter is run to obtain updated position and attitude parameters. After this first update, the filter is run after every subsequent image to continually adjust the state parameters. These new



parameters are then used by the attitude control system to point the camera at the right spot for the remainder of the image opportunities.

There are two reasons why the filter is not immediately invoked after the first image. The first is to protect against any image, especially the first, having a bad centroid and corrupting the solution. The second is to maintain open loop tracking using the best ground information as long as possible so that if the tracking failed, there is a reasonable chance that some science images would still be available. It is assumed that after the first update, the solution is fairly stable and not as easily corrupted by bad data. Then, the updates are performed at every image to keep up with the rapidly changing geometry which provides the best information on the downtrack position of the spacecraft.

## Testing and Verification

### Covariance Results

The first step in verifying that the input parameters, data weights, and image timing are sufficient to maintain lock on the target is to examine the covariances output from the filter. The most useful way to visualize the covariance is to rotate into the Radial-Transverse-Normal (*RTN*) coordinate system. In this system,  $R$  is the vector from the spacecraft to the target,  $N$  is the cross product of  $R$  with the velocity vector, and  $T$  completes the right-handed triad. The significance of using this coordinate system is that optical data taken in the months prior to encounter provide fairly good information about  $T$  and  $N$  which are perpendicular to the incoming asymptote, but  $R$ , which corresponds to the time to encounter, is less well known. Furthermore, while knowledge of  $T$  and  $N$  improves immediately and is well determined prior to encounter,  $R$  depends on the changing parallax to reduce its uncertainty; depending on the flyby distance, this change may happen fairly late.

An example of the filtered position uncertainties as a function of time are shown in Figure 2. This example is specific for the STARDUST Wild 2 encounter. The *a priori* position covariance was set at 1100 km for  $R$  and 20 km for  $T$  and  $N$ ; the *a priori* attitude uncertainty was set to 0.1 deg., spherical. Imaging was started at Encounter (E) minus 30 minutes, with images taken at 30 second intervals. The filter was updated at E-10 minutes, and starting at E-6 minutes, the image frequency increased to one every 20 seconds. Given the planned flyby distance of 250 km and the camera FOV, the goal is to know the spacecraft trajectory to 6.5 km or better; any worse and the comet would drift out of the camera FOV. The left panel of Figure 2 shows the radial uncertainty (for ease of visualization, the plot is logarithmic in  $y$ ), with the horizontal dashed line being the 6.5 km goal. Note that this level of knowledge comes very late, stressing the need to keep tracking all the way through. The right panel of Figure 2 shows the uncertainties of the other two position components; of these, the  $N$  direction is critical because this determines the magnitude of the roll needed to align the scan mirror plane with the flyby plane. The plot shows the 6.5 value is reached at around seven minutes prior to encounter, sufficient to compute the correct roll.

### Monte Carlo Results

If the dynamic equations used in the filter precisely modeled the true forces acting on the spacecraft, then the covariance obtained after filtering would accurately represent the statistics of the estimated values. This is clearly not the case however, as we have deliberately used a reduced set of dynamics to keep the algorithm simple and fast. For this reason, Monte Carlo simulations are needed to assess the ability of the algorithm to maintain visual lock on the nucleus. For the simulations, a “truth” model of the trajectory, spacecraft attitude, and observations are generated and provided to the filter. For a given run, the truth model represents a random sampling of the error sources which affect that model. One-hundred runs are performed, and the results are evaluated by determining whether or not the nucleus was visible in the camera FOV at all times.

To incorporate realism into the simulations, the truth trajectory is propagated using two-body Keplerian motion rather than the straight line used in RSEN. The spacecraft attitude knowledge

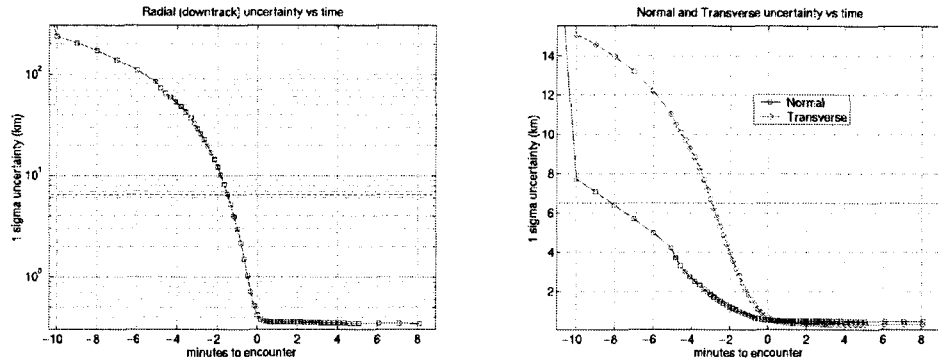


Figure 2: Post-fit Position Sigmas

is sampled using realistic models of the gyro bias and drift behavior, so the end result mimics fairly accurately what the attitude errors in flight would be. Images are created using a simulated, irregularly shaped object with correct lighting conditions, and each sample of the Monte Carlo run uses a different orientation for the object. The sampled error sources therefore include errors in the initial state and attitude, gyro bias and drift parameters, and orientation and rotation of the target. All errors are sampled from a zero mean, Gaussian distribution with appropriate sigmas.

The ultimate criteria for success is simply whether or not the target is maintained in the camera FOV. In particular for STARDUST, the goal was to obtain at least 30 images of the comet within a range of 2000 km, and obtain at least one image at a distance of less than 300 km. To distill the results from the simulations, several plots were generated from each run, two of which are shown here. The first is in Figure 3. Here, the x-axis is the time to encounter; the y-axis plots, for each image in the sequence, the number of frames out of the 100 samples which had the comet in the FOV (the line with the squares). Thus, a value of 100 for a particular time means that all 100 samples at that time had the comet. The plot shows that for the first 15 minutes, the success rate was 100% which then drops to 80% at E-10 minutes. Then, the filter update happens which bumps up the success to 98%. A smaller drop is then seen until the roll executes at E-6 minutes, and the remaining are all successful.

Since the important images are the ones at distances of 2000 km though, the other line in Figure 3 (the circles) shows the number of frames which include the target and is also at less than 2000 km range. Thus, even though the early images are all in the FOV, the distance is greater than 2000 km, so the probability of meeting the success criteria is 0. After about E-10 minutes, an increasing number of cases does meet the criteria. The peak values near the nominal closest approach reach 99% probability of success.

A second plot, Figure 4, shows a histogram of the spacecraft to comet range at the closest successful image over the 100 samples. The peak is at the nominal flyby distance of 250 km, with 50% of the samples at this range. The solid line is the cumulative probability of the range at the closet image; although not visible on this scale, this indicates an overall probability of 98% of the samples meeting the minimum 300 km distance criteria.

Similar plots were also generated for the Borrelly and Annefrank encounters. These were a useful tool in tweaking the filter and other parameters to optimize performance, and increased confidence that RSEN would be successful in tracking through the flybys.




Figure 3: Success rate for each image opportunity during encounter.




Figure 4: Histogram of distance when closest image was shuttered.

## Flight Results

### Deep Space 1 at Borrelly

The DS1 flyby of comet Borrelly occurred on September 22, 2001 at 22:30:33 ET. The flyby distance at closest approach was roughly 2100 km, with a velocity of 16.6 km/s. RSEN was initiated at about E-32 minutes; the ground-based navigation information used to initialize the state had optical data up to E-12 hours.<sup>6</sup> The cross-track accuracy of the ground navigation solution was about 15 km, but to keep the RSEN filter from being too constrained, the *a priori* sigma on the state was kept at 20 km. The radial (downtrack) direction of the spacecraft to comet was largely determined from ground-based observations of the comet which took place all through the approach phase; the position was known to better than 200 km, but for RSEN, the sigma was once again set to a slightly larger value of 350 km. The *a priori* sigmas on the gyro bias and drift was set to 150 pixels and 400 pix/hr, respectively, for both the pixel and line directions; this corresponds to 0.1 deg and 0.3 deg/hr. Images were then shuttered at a rate of approximately one every 30 seconds, with exposures ranging from 76 to 600 msec (the large range was to account for uncertainties in the brightness of Borrelly). A total of 52 images was taken for RSEN, and the first RSEN updated solution used for tracking occurred at E-10 minutes.

Figure 5 schematically plots a track of the observed brightness centroids as determined by RSEN in the camera FOV (selected actual images are shown in Figure 6). At the start, Borrelly was roughly centered at around pixel and line 600. As time went on, Borrelly remained more or less centered until about E-11 minutes. At this time, there was a gap in the RSEN images to allow a second instrument, the Infrared (IR) Detector, to scan across Borrelly, which took Borrelly out of the camera FOV. After the IR scan, Borrelly was located near the bottom of the frame, but RSEN slowly brought it back to near the center. About E-5 minutes, the ACS attitude mode was updated; subsequently, the images stayed centered in the pixel direction but moved to near the top of the frame in line. The reason for this shift has not been determined. Examination of the state and gyro solutions from RSEN indicate that the solution had converged at this point and the residuals showed no sudden change. From this, it can be inferred that the predicted value of Borrelly matched the observed value, so the RSEN estimated solution was correct. For reasons never conclusively determined, the ACS controller on the spacecraft placed the boresight slightly lower than needed. Nevertheless, all expected images of Borrelly were obtained, with the closest image shuttered at E-2 min, 46 seconds, at a distance of 3514 km and a resolution of 46 m per pixel.

### STARDUST at Annefrank

The STARDUST flyby of asteroid Annefrank occurred on November 2, 2002 at 04:51:19 ET. The flyby distance at closest approach was 3076 km, with a velocity of 7.23 km/s. The primary purpose of the Annefrank encounter was to perform as complete an engineering test of all flyby events as a preparation for the Wild 2 encounter. The test of RSEN was especially important as it was the only opportunity to exercise the STARDUST version of RSEN in flight. The large flyby distance was chosen to minimize any likelihood of impact with the asteroid; as a result, the extent of the Annefrank images, even at closest approach, was less than 30 pixels.

RSEN was initiated at E-20 minutes with a target-relative state based on radio data alone. Optical navigation images of Annefrank were planned and taken between E-38 hours and E-12 hours, but the extreme geometry of the encounter (the phase angle on approach was 150 deg), coupled with the small size of Annefrank (3 km) resulted in Annefrank being too dim to image at these times. Thus, the *a priori* state initialization of RSEN relied on ground-based ephemeris information for the asteroid, combined with the radio-based estimate of the spacecraft's position and velocity. Although the predicted error of Annefrank's ephemeris was better than 80 km (1 sigma), RSEN was initialized with a more conservative value of 100 km for the *a priori* uncertainty in all three (*RTN*) components. The *a priori* uncertainty for the gyro bias was set to 0.1 deg (1 sigma) in all three axes. Images were taken at approximately 30 second intervals, and the filter update




Figure 5: Observed brightness centers of Borrelly through encounter.




Figure 6: RSEN images of borrelly through encounter.

occurred at E-10 minutes. The time of the roll maneuver needed to align the mirror plane with the flyby plane was set at E-4 minutes. RSEN was terminated at E+3 minutes.

Figure 7 plots the observed brightness centroid track in the FOV. Note that the images start very near the upper left corner of the frame, and actually left the frame at E-11 minutes. This was largely due to the error in the *a priori* downtrack spacecraft position. After the filter update at E-10 minutes, however, the asteroid was brought back to the center of the FOV, and it remained very well centered through closest approach and after. Interestingly, although the roll angle was computed, it was never executed due to the fact that the *a priori* out-of-plane position error was less than the threshold roll angle of 0.3 deg. The result was entirely fortuitous since the lack of optical navigation frames of Annefrank during the approach meant no target-relative information was available.




Figure 7: Observed brightness centers of Annefrank through encounter.

Figure 8 plots the *RTN* position corrections computed by RSEN at each of the state update times. Note that the first correction made at E-10 minutes changed very little over the course of the encounter. The final corrections made by RSEN were -89 km in *R* (or equivalently, 12 seconds in the time of closest approach), -174 km in *T*, and -8 km in *N*. The values in *R* and *T* were quite a bit larger than the expected errors, but RSEN did not have trouble with the corrections. The closest image was shuttered less than a second before closest approach at a distance of 3079 km and a resolution of 185 m. The fact that the asteroid was tracked both on the inbound and outbound asymptotes provided enough information to compute its phase curve, an important science result which would not have been available using mosaics<sup>7</sup>.

## STARDUST at Wild 2

The STARDUST encounter with comet Wild 2 took place on January 2, 2004 at 19:22:36 ET. The flyby distance at closest approach was 237 km at a velocity of 6.12 km/s. Although the main purpose of the flyby was to collect coma dust samples, images of the comet nucleus were highly desired as an additional science goal. The flyby distance was largely chosen to maximize the likelihood of collecting the required amount of dust while minimizing the chance of damaging the spacecraft. The geometry of the flyby was such that the phase angle on approach was 72 deg, and at closest approach, it was about 11 deg, nearly fully illuminating the comet.

RSEN was initialized at E-30 minutes with a comet relative state based on combined radio data of the spacecraft with optical images of the comet. The latter was especially important since the comet




Figure 8: RSEN filtered position corrections in the  $RTN$  coordinate system for Annefrank.

was behind the sun as viewed from Earth through much of the approach phase. The *a priori* comet ephemeris was thus based on ground-based astrometry from late May 2003, and the uncertainty in the comet's location projected to encounter was nearly 2000 km (1 sigma). With the final approach opnav images taken at E-14 hours, the spacecraft's comet relative position accuracy was about 5 km in the crosstrack directions ( $T$  and  $N$ ). The downtrack ( $R$ ) could not be improved by the optical navigation frames, however, the geometry was such that this direction was directly observable from the Earth. An intensive ground observing campaign was initiated as soon as the comet appeared from behind the sun in late December. This campaign resulted in the ephemeris uncertainty being reduced from 1100 km to 300 km (1 sigma) in the downtrack direction. Nevertheless, based on results from the Monte Carlo simulations described earlier, RSEN was initialized with a covariance of 1100x20x20 km in the  $RTN$  frame. The gyro bias was once again initialized with *a priori* uncertainties of 0.1 deg in all three axes.

Images were shuttered with a frequency of once every 30 seconds from E-30 to E-6 minutes, then once per 10 seconds from E-6 to E+6 seconds, and then again every 30 seconds to E+8 minutes. Only every other of the 10 second images were used by RSEN; the remaining were solely for science. The filter update occurred at E-10 minutes, and the roll was performed at E-6 minutes. Figure 9 plots the observed brightness centroids through the encounter. The comet started out well centered and slowly drifted downwards. Before the filter update, the comet was located around line 650; after the update, it shifted back to near where it started at near line 475. Another shift, this time in the pixel direction, occurred after the roll maneuver, whose magnitude was about 4 deg. A further rapid change in pixel occurred as the spacecraft passed through closest approach where state errors had the largest mapping in the camera FOV. Overall, however, the comet stayed well within the FOV through the entire encounter.

Figure 10 plots the  $RTN$  position corrections computed by RSEN as a function of time. The final converged state update was 530, -20, and -11 km in  $RTN$ , respectively. In  $T$  and  $N$ , the initial update at E-10 minutes was very near the final; for  $R$ , the convergence only occurred at about E-5 minutes. The correction in  $R$  amounted to a time-of-flight change of over 86 seconds which was somewhat larger than expected, but still within the ability of RSEN to handle. The closest image shuttered was 3 seconds prior to encounter at a distance of 237 km and a resolution of 14 m. This and the other frames surrounding closest approach are the highest resolution images of a comet to date.




Figure 9: Observed brightness centers of Wild 2 through encounter.




Figure 10: RSEN filtered position corrections in the *RTN* coordinate system for Wild 2.



## CONCLUSIONS

Flybys of small solar system bodies provide much useful science information about the nature of these objects, including size, shape, and composition. To maximize science return, it is important to have the best available knowledge of the trajectory of the spacecraft relative to the object, both to point the camera and any other instruments. Since the best information about the trajectory is provided by close encounter images, a methodology to use this data autonomously onboard the spacecraft is very desirable. A very robust system to track a small body during the flyby using the near encounter images was developed, tested, and proven in flight. The results show the advantages in science return when using the tracking system as opposed to following a mosaic pattern, which, although successful, results in unusable image frames. The robustness of the approach described has been proven by the success of the tracking system on three separate occasions under different lighting, flyby velocity, and other geometric conditions.

## ACKNOWLEDGEMENTS

The work described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

## REFERENCES

1. S. Bhaskaran, J. E. Riedel, and S. P. Synnott, "Autonomous Nucleus Tracking for Comet/Asteroid Encounters: The STARDUST Example", Paper 97-628, AAS/AIAA Astrodynamics Specialist Conference, Sun Valley, Idaho, August 4-7, 1997
2. W. M. Owen and R. M. Vaughan, "Optical Navigation Program Mathematical Models", JPL Internal Document JPL-EM 314-513, August 9, 1991.
3. S. P. Synnott, A. J. Donegan, J. E. Riedel, J. A. Stuve, "Interplanetary Optical Navigation: Voyager Uranus Encounter", Paper 86-2113, AIAA/AAS Astrodynamics Conference, Williamsburg, Virginia, August 1986.
4. J.C. Russ, *The Image Processing Handbook*, CRC Press, IEEE Press, Chapter 7, 1999.
5. R. A. Werner, personal communication.
6. S. Bhaskaran, J. E. Riedel, B. M. Kennedy, T. C. Wang, "Navigation of the Deep Space 1 Spacecraft at Borrelly", Paper AIAA 2002-4815, AIAA/AAS Astrodynamics Specialist Conference, Monterey, California, August 5-8, 2002.
7. R. L. Newburn et al., "Phase curve and albedo of asteroid 5535 Annefrank", *J. of Geophysical Research*, 108(E11), 5117, doi:10.1029/2003JE002106, 2003.